IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Alfredo Aldereguia, *et al.*	§	Group Art Unit: 2619
	§	
	§	Examiner: Chu, Wutchung
Serial No.: 10/777,508	§	
	§	Atty Docket No.: RPS920030207US1
Filed: 02/12/2004	§	
	§	Customer No.: 56102
Title: Automated Topology Detection In A Data Processing System	§	
	§	Confirmation No.: 5244

Mail Stop: Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

APPEAL BRIEF

Honorable Commissioner:

This is an Appeal Brief filed pursuant to 37 CFR § 41.37 in response to the Final Office Action of April 16, 2008 (hereinafter the "Office Action"), and pursuant to the Notice of Appeal filed July 16, 2008.

REAL PARTY IN INTEREST

The real party in interest in accordance with 37 CFR § 41.37(c)(1)(i) is the patent assignee, International Business Machines Corporation ("IBM"), a New York corporation having a place of business at Armonk, New York 10504.

RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences within the meaning of 37 CFR § 41.37(c)(1)(ii).

STATUS OF CLAIMS

Status of claims in accordance with 37 CFR § 41.37(c)(1)(iii): Twenty-one (21) claims are filed in the original application in this case. Claims 1-21 are rejected in the Office Action. Claims 1-21 are on appeal.

STATUS OF AMENDMENTS

Status of amendments in accordance with 37 CFR § 41.37(c)(1)(iv): No amendments were submitted after final rejection. The claims as currently presented are included in the Appendix of Claims that accompanies this Appeal Brief.

SUMMARY OF CLAIMED SUBJECT MATTER

Appellants provide the following concise summary of the claimed subject matter according to 37 CFR § 41.37(c)(1)(v). This summary includes a concise explanation of the subject matter defined in each of the independent claims involved in the appeal and includes references to the specification by page and line number and to the drawings by elements. The three independent claims involved in this appeal are claims 1, 8, and 15. Claim 1 is a system claim. Claim 8 is a method claim. Claim 15 recites computer program aspects of the method of claim 8.

Claim 1 recites a scalable data processing system. The scalable data processing system of claim 1 includes a first set of central processing units (page 4, lines 8-18; Figure 3, elements 101-104). The scalable data processing system of claim 1 also includes a first system memory accessible to the first set of processors (page 4, lines 8-18; Figure 3, elements 101-104 and 135). The scalable data processing system of claim 1 also includes scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system (page 5, line 20 – page 6, line 29; Figure 3, elements 100-1, 100-2, 135, and 101-104). The scalable data processing system of claim 1 also includes a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first

system to the second system (page 6, line 29 – page 7, line 21; Figure 3, elements 100-1, 100-2, 151-153, and 202). The scalable data processing system of claim 1 also includes system management to cause each of the system's scalability ports to issue an identifiable signal and further configured to detect the reception of an identifiable signal, sent by another system, by any of the scalability ports and to report the reception of the signal to a system management of the second system to determine which ports of the two systems are connected by the cable (page 6, line 29 – page 7, line 21; Figure 3, elements 100-2, 100-1, 151-153, and 320).

Claim 8 recites a method of determining scalability cabling between at least two scalable data processing systems (page 7, lines 1-8; Figure 3, elements 100-1 and 100-2). The method of claim 8 includes driving an identifiable signal on a first scalability port of a first system (page 7, lines 8-10; Figure 3, elements 151, 100-1, and 320). The method of claim 8 also includes responsive to receiving the identifiable signal by a second system, determining which scalability port of the second system received the distinctive signal (page 7, lines 14-21; Figure 3, element 320). The method of claim 8 also includes informing the first system of the reception of the distinctive signal by the determined scalability port of the second system and recording the first scalability port of the first system and the scalability port of the second system as being connected by a scalability cable (page 7, lines 22-28; Figure 3, elements 312, 100-1, and 100-2).

Claim 15 recites a computer program product for determining scalability cabling between at least two scalable data processing systems (page 7, lines 1-8; Figure 3, elements 100-1 and 100-2). The computer program product of claim 15 includes computer executable instructions stored on a computer readable medium (page 8, lines 5-8). The computer program product of claim 15 also includes computer code means for driving an identifiable signal on a first scalability port of a first system (page 7, lines 8-10; Figure 3, elements 151, 100-1, and 320). The computer program product of claim 15 also includes responsive to receiving the identifiable signal by a second system, computer code means for determining which scalability port of the second system received the distinctive signal (page 7, lines 14-21; Figure 3, element 320). The computer program product of claim 15

also includes computer code means for informing the first system of the reception of the distinctive signal by the determined scalability port of the second system and recording the first scalability port of the first system and the scalability port of the second system as being connected by a scalability cable (page 7, lines 22-28; Figure 3, elements 312, 100-1, and 100-2).

GROUNDS OF REJECTION

In accordance with 37 CFR § 41.37(c)(1)(vi), Appellants provide the following concise statement for each ground of rejection:

1. Claims 1-2, 8-10, and 15-17 stand rejected under 35 U.S.C. § 102(b) over Leigh *et al.*, (U.S. Publication No. 2003/0158940).

2. Claims 3, 6-7, 11-12, 14, 18-19, and 21 stand rejected for obviousness under 35 U.S.C. § 103(a) as being unpatentable over Leigh in view of Nakamura (U.S. Publication No. 2004/0057448).

3. Claim 4 stands rejected for obviousness under 35 U.S.C. § 103(a) as being unpatentable over Leigh in view of Nakamura further in view of Vegter (U.S. Patent 6,286,073).

4. Claims 5, 13, and 20 stand rejected for obviousness under 35 U.S.C. § 103(a) as being unpatentable over Leigh in view of Vegter.

ARGUMENT

Appellants present the following argument pursuant to 37 CFR § 41.37(c)(1)(vii) regarding the ground of rejection on appeal in the present case.

Argument Regarding The First Ground Of Rejection On Appeal: Claims 1-2, 8-10, And 15-17 Are Rejected Under 35 U.S.C. §102(b) Over Leigh

Claims 1-2, 8-10, and 15-17 stand rejected under 35 U.S.C. § 102 as being anticipated by Leigh, *et al.* (U.S. Publication No. 2003/0158940) (hereafter, 'Leigh'). To anticipate claims 1-2, 8-10, and 15-17 under 35 U.S.C. § 102, Leigh must disclose and enable each and every element and limitation recited in the claims of the present application. As explained below, Leigh does not disclose and enable each and every element and limitation recited in the claims of the present application and therefore does not anticipate the claims of the present application.

Leigh Does Not Disclose Each and Every Element Of Claim 1 Of The Present Application

"A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Independent claim 1 recites:

> 1. A scalable data processing system, including:
>
> a first set of central processing units;
>
> a first system memory accessible to the first set of processors;
>
> scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system;
>
> a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system; and

> system management to cause each of the system's scalability ports to issue an identifiable signal and further configured to detect the reception of an identifiable signal, sent by another system, by any of the scalability ports and to report the reception of the signal to a system management of the second system to determine which ports of the two systems are connected by the cable.

As explained in more detail below, Leigh does not disclose each and every element of claim 1, and Leigh therefore cannot be said to anticipate the claims of the present application within the meaning of 35 U.S.C. § 102.

Leigh Does Not Disclose Scalability Logic To Connect The Data Processing System To A Second Data Processing System, Having A Second Set Of Processors And A Second System Memory, To Form A Scaled System

The Office Action takes the position that Leigh at paragraph 0029 and figure 1c discloses scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system. Appellants respectfully note in response, however, that what Leigh at paragraph 0029, in fact discloses is:

> [0029] The four-ILB embodiment shown in FIG. 1(c) can be used as an example to describe the zone-based load balancing method. Assume that ILBs 10 and 20 belong to zone-1 and ILBs 30 and 50 belong to zone-2 and that ILBs 10 and 50 are the primary master ILBs for their respective zones. Simpler load-balancing algorithms, such as a basic round-robin method, may be used when the slave ILBs do not overlap across the zones. In the example of FIG. 1(c), slave ILB 20 is in zone-1 only and slave ILB 30 is in zone-2 only, although it is possible to assign ILB 20 and ILB 30 to be in both zone-1 and zone-2. A primary master ILB, e.g., ILB 10, can collaborate with another primary master ILB, e.g., ILB 50, for load sharing. Each primary master ILB has its own zone as its primary zone and the collaborating zone as its secondary zone. A primary master ILB may decide to cross the zone boundaries and send some of its load to the servers in another zone based on based on factors such as its host server's

> workload index, network data traffic level, server health, and ILB health. "Server health" constitutes functional statuses on the server host's critical subsystem components, such as CPU(s), cache memory, system memory, and disks. "ILB health" constitutes functional statuses on the ILB's components, such as processors, memory buffers, ASICs, and FPGAs.

That is, Leigh at paragraph 0029, discloses zone-based load balancing implemented with Internal Load Balancers or 'ILBs.' Leigh's ILBs are special purpose integrated circuits such as field programmable gate arrays ('FPGAs') or application specific integrated circuits ('ASICs'). Leigh's load balancing with ILBs does not disclose scalability logic connecting one data processing system to a second data processing system as claimed here. It is clear to a person of skill in the art that load balancing as disclosed in Leigh simply has nothing to do with scalability logic for connections among systems. In Leigh, the servers among whom load balancing is implemented are expressly already connected in a network, therefore requiring no scalability logic for connections, as claimed in the present application. Indeed, there is only one substantive mention of scaling in Leigh, in paragraph 0057, addressing only sleeping and waking among servers already connected in a network. On the face of Leigh, therefore, there would be no reason for a person of skill in the art to expect to derive anything regarding scalability for connections among systems from Leigh, which is clearly about load balancing and not about scalability.

In addition to the fact that Leigh generally discloses nothing about scalability logic connecting data processing systems, there is another more specific reason why Leigh does not disclose scalability logic connecting one data processing system to a second data processing system: Leigh does not disclose a 'scaled system' as claimed in the present application. In the original specification at page 4, lines 4-7, Appellants define a 'scaled system' as a system in which two or more symmetric multiprocessor ("SMP") systems are interconnected to form a larger multiprocessor system capable of functioning under a single operating system image. Neither at the reference point cited by the examiner nor anywhere else in Leigh does Leigh disclose a scaled system capable of functioning under a single operating system image. In fact, the terms "image," "operating system," and "operating system image" do not appear at any point in Leigh – not even once. For these reasons, a person of skill in the art would understand a 'scaled system,' as the term in

used in the present claims read in view of the specification, to be interconnected SMPs capable of running under a single operating system image. Moreover, it is clear to a person of skill in the art that all the servers described in Leigh are assumed to be operating their own separate operating systems. Because Leigh does not disclose a 'scaled system' as claimed in the present application, Leigh cannot possibly disclose scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a 'scaled system' as claimed here.

In addition to the fact that Leigh does not disclose a scaled system, there is another reason that Leigh does not disclose scalability logic connecting one data processing system to a second data processing system: Leigh does not disclose the 'data processing system' claimed in the present application. In the original specification at page 4, lines 11-12, Appellants define a 'data processing system' as including a Processor Scalability and Cache control unit which controls access to the scalability ports in the data processing system. A person of skill in the art therefore would understand a 'data processing system,' as the term is used in the present claims read in view of the specification, to include a Processor Scalability and Cache control unit which controls access to scalability ports. Leigh does not disclose a data processing system containing scalability ports or a data processing system containing a Processor Scalability and Cache control unit which controls access to the scalability ports. Because Leigh does not disclose a data processing system as claimed in the present application, Leigh cannot possibly disclose scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system. Because Leigh does not disclose each and every element and limitation of Appellants' claims, Leigh does not anticipate Appellants' claims, and the rejections under 35 U.S.C. § 102 should be withdrawn.

Leigh Does Not Disclose A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System

The Office Action takes the position that Leigh at paragraphs 0029, 0040, and figures 1a-1d, discloses a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system. Appellants respectfully note in response, however, that as discussed above, Leigh at paragraph 0029 in fact discloses zone-based load balancing implemented with ILBs. In addition, what Leigh at paragraph 0040 in fact discloses is:

> [0040] The use of 3-port ILB modules allows the creation of load-balancing clusters at low cost, since no additional hardware, such as a network switch, is needed. This is important especially when there are only a few servers to be load balanced and an ILB module in each host server can provide the load-balanced cluster solution. When there are only a few servers, the installation of the interconnecting cables among the ILBs is not a significant issue. FIG. 1(d) shows an example of how five ILBs 10, 20, 30, 40, and 50 can be interconnected among themselves and to two external networks via network segments 80 and 90. Six inter-ILB connections 83, 84, 85, 86, 87, and 88 are used to interconnect the five 3-port ILBs. The configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method. Similarly, in the embodiment shown in FIG. 2, twelve 3-port ILBs 202 through 224 are interconnected by sixteen connections and four primary master ILBs 202, 208, 214, and 220 are connected to four external networks via connections 230, 240, 250, and 260. The number of possible permutations of connections among the ports and the number of sub-optimal configurations are far greater in this embodiment. In either case, a technician installing the ILBs may have difficulty in configuring the connections in the most efficient manner for the load balancing method to be used by the ILBs. Therefore, a standard non-load-balancing switch may be used to expedite the interconnection process when several ILBs, e.g., five or more, are present.

That is, Leigh at paragraph 0040 discloses permutations of arrangements of connections among ports of ILBs. Neither Leigh's disclosure of zone-based load balancing implemented with ILBs nor Leigh's disclosure of permutations of arrangements of

connections among ports of ILBs disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system as claimed in the present application. Leigh does not disclose scalability cables, scalability ports, or scalability logic. Scalability cables are distinct from network cables, such as Ethernet cables, that are used for connecting machines within a local area network. *See*, Appellants' original specification at page 1, lines 16-18. Neither at the reference point cited by the Examiner nor anywhere else in Leigh does Leigh disclose scalability cables. Leigh merely discloses that servers are connected through the network ports of Leigh's internal load balancers, without making any disclosure regarding the cables that are used to connect Leigh's servers. Because Leigh does not disclose anything regarding the cables used to connect Leigh's servers, Leigh cannot disclose scalability cables.

In addition to the fact that Leigh does not disclose scalability cables, there is another reason that Leigh does not disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system: Leigh does not disclose scalability ports. Scalability ports, as claimed in the present application, receive scalability cables connecting two systems. As discussed above, Leigh does not disclose scalability cables. Because Leigh does not disclose scalability cables, Leigh cannot possibly disclose scalability ports that receive scalability cables, as claimed here. Because Leigh does not disclose scalability cables or scalability ports, Leigh cannot disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system as claimed in the present application.

In addition to the fact that Leigh does not disclose scalability cables or scalability ports, there is another reason that Leigh does not disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system: Leigh does not disclose scalability logic. Scalability logic, as claimed in the present application, is connected to scalability ports. As discussed above, Leigh does not disclose scalability ports. Because Leigh does not disclose scalability ports, Leigh cannot

possibly disclose scalability logic connected to scalability ports, as claimed here. Because Leigh does not disclose scalability cables, scalability ports, or scalability logic, Leigh cannot disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system as claimed in the present application. Because Leigh does not disclose each and every element and limitation of Appellants' claims, Leigh does not anticipate Appellants' claims, and the rejections under 35 U.S.C. § 102 should be withdrawn.

Leigh Does Not Disclose System Management To Cause Each Of The System's Scalability Ports To Issue An Identifiable Signal And Further Configured To Detect The Reception Of An Identifiable Signal As Claimed In The Present Application

The Office Action takes the position that Leigh at paragraphs 40, 30, and 47, and Figures 1a -1d, discloses: system management to cause each of the system's scalability ports to issue an identifiable signal and further configured to detect the reception of an identifiable signal, sent by another system, by any of the scalability ports and to report the reception of the signal to a system management of the second system to determine which ports of the two systems are connected by the cable. Appellants respectfully note in response, however, that what Leigh at paragraph 40 in fact discloses is permutations of arrangements of connections among ports of ILBs. In addition, what Leigh at paragraph 30, in fact discloses is:

> [0030] As an example of zone boundary crossing, in FIG. 1(c), after primary master ILB 10 has distributed the incoming load from network segment 80 to a predefined saturation level on the host servers associated with ILB 10 and ILB 20 in zone-1, ILB 10 may interrogate the other primary master, ILB 50 in zone-2, to possibly accept future incoming loads on network segment 80. If ILB 50 acknowledges the load shedding request from ILB 10, then ILB 50 provides ILB 10 with a list of ILBs that can accept the load and with the load-shedding conditions. The load-shedding conditions may include information such as absolute time interval, load-shed check timer value, and number of loads. An example of "absolute time interval" is a wall-clock time interval preset by system administrators based on the known load condition, such as during the first two hours of every weekday, or every business day at lunch hours. An

> example of "load-shed check timer value" is a number of hours or minutes, chosen by system administrators, to be set in a register and counted down. This timer value can be set at the time when a load reaches a predefined load index threshold or it can be set periodically. When the timer value has counted down, the corresponding ILB will check the load condition with respect to its resource capacity to determine whether or not it should notify the primary master ILB (if it is not one itself) to redirect the load to another ILB. If ILB 10 sheds its load while complying with the load-shedding conditions, it will stop the load-shedding activity upon a load-shed-abort signal from ILB 50. ILB 50 may issue this load-shed-abort signal to ILB 10 when the servers whose loads are being shed fall below a predefined resource saturation threshold.

And what Leigh at paragraph 0047 in fact discloses is:

> [0047] At the beginning of the topology discovery algorithm, each ILB will command its individual ports to broadcast a data packet known as a roll-call-1-query packet. Each ILB port will respond to the roll-call-1-query packet with a roll-call-1-response packet that includes information such as its port address, the associated ILB ID, and the associated host server ID. Since this query packet is comprehensible only by a port on an ILB, a port will know whether it is connected to another ILB port, to an external network, or to nothing at all. After this step, each port knows the ID of the port or ports to which it is directly connected. For the topologies shown in FIGS. 1(a) through 1(d) where the ILBs are interconnected directly to each other rather than through a switch, each ILB port knows the other ILB port it is connected to, if any. For example, in FIG. 1(b), ports 12 and 56 know that they are connected to external networks and ports 14, 16, 22, 26, 52, and 54 know that they are connected to the corresponding ports on the other ILBs, i.e., to ports 54, 22, 16, 52, 26, 14, respectively. Port 24 knows it is not connected.

That is, Leigh at paragraphs 30 and 47 discloses a topology detection algorithm that broadcasts a packet and receives a responsive packet. Neither Leigh's permutations of arrangements of connections among ports of ILBs nor Leigh's topology detection algorithm that broadcasts a packet and receives a responsive packet disclose system management with scalability ports, as claimed in the present application. In particular, Leigh does not disclose scalability ports. An identifiable signal, according to the claims in the present application, is broadcast and received through the scalability port of a system. Without disclosing a scalability port, Leigh cannot disclose system management

using scalability ports as claimed in the present application. Because Leigh does not disclose each and every element and limitation of Appellants' claims, Leigh does not anticipate Appellants' claims, and the rejections under 35 U.S.C. § 102 should be withdrawn.

Leigh Does Not Disclose Each and Every Element Of Claim 8 Of The Present Application

"A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Independent claim 8 recites:

> 8. A method of determining scalability cabling between at least two scalable data processing systems, comprising:
>
> driving an identifiable signal on a first scalability port of a first system;
>
> responsive to receiving the identifiable signal by a second system, determining which scalability port of the second system received the distinctive signal;
>
> informing the first system of the reception of the distinctive signal by the determined scalability port of the second system and recording the first scalability port of the first system and the scalability port of the second system as being connected by a scalability cable.

As explained in more detail below, Leigh does not disclose each and every element of claim 8, and Leigh therefore cannot be said to anticipate the claims of the present application within the meaning of 35 U.S.C. § 102.

Leigh Does Not Disclose Driving An Identifiable Signal On A First Scalability Port Of A First System

The Office Action takes the position that Leigh at paragraphs 0040 and 0030 discloses: driving an identifiable signal on a first scalability port of a first system. Appellants respectfully note in response, however, that what Leigh at paragraph 40, in fact discloses is permutations of arrangements of connections among ports of ILBs. In addition, what Leigh at paragraph 0030, in fact discloses is:

> [0030] As an example of zone boundary crossing, in FIG. 1(c), after primary master ILB 10 has distributed the incoming load from network segment 80 to a predefined saturation level on the host servers associated with ILB 10 and ILB 20 in zone-1, ILB 10 may interrogate the other primary master, ILB 50 in zone-2, to possibly accept future incoming loads on network segment 80. If ILB 50 acknowledges the load shedding request from ILB 10, then ILB 50 provides ILB 10 with a list of ILBs that can accept the load and with the load-shedding conditions. The load-shedding conditions may include information such as absolute time interval, load-shed check timer value, and number of loads. An example of "absolute time interval" is a wall-clock time interval preset by system administrators based on the known load condition, such as during the first two hours of every weekday, or every business day at lunch hours. An example of "load-shed check timer value" is a number of hours or minutes, chosen by system administrators, to be set in a register and counted down. This timer value can be set at the time when a load reaches a predefined load index threshold or it can be set periodically. When the timer value has counted down, the corresponding ILB will check the load condition with respect to its resource capacity to determine whether or not it should notify the primary master ILB (if it is not one itself) to redirect the load to another ILB. If ILB 10 sheds its load while complying with the load-shedding conditions, it will stop the load-shedding activity upon a load-shed-abort signal from ILB 50. ILB 50 may issue this load-shed-abort signal to ILB 10 when the servers whose loads are being shed fall below a predefined resource saturation threshold.

That is, Leigh at paragraph 0030, discloses distributing network loads across servers that are in different zones of a network. Neither Leigh's permutations of arrangements of connections among ports of ILBs nor Leigh's distributing network loads across servers that are in different zones of a network discloses driving an identifiable signal on a first

scalability port of a first system because Leigh does not disclose scalability ports. In the original specification at page 2, line 5-7, Appellants define scalability ports as receiving scalability cables. As discussed above, Leigh does not disclose scalability cables. Because Leigh does not disclose scalability cables, Leigh cannot possibly disclose scalability ports that receive scalability cables, as claimed in the present application. Because Leigh does not disclose scalability ports as claimed in the present application, Leigh cannot possibly disclose driving an identifiable signal on a first scalability port of a first system as claimed here.

In addition to the fact that Leigh does not disclose driving an identifiable signal on a first scalability port of a first system because Leigh does not disclose scalability ports as claimed here, for identical reasons, Leigh also does not disclose the remaining elements of claim 8. Each remaining element of claim 8 recites a scalability port. Therefore, Leigh cannot possibly disclose the remaining elements of claim 8 because Leigh does not disclose a scalability port as claimed here. Because Leigh does not disclose each and every element and limitation of Appellants' claims, Leigh does not anticipate Appellants' claims, and the rejections under 35 U.S.C. § 102 should be withdrawn.

Leigh Does Not Enable Each and Every Element Of The Claims Of The Present Application

Not only must Leigh disclose each and every element of the claims of the present application within the meaning of *Verdegaal* in order to anticipate Appellants' claims, but also Leigh must be an enabling disclosure of each and every element of the claims of the present application within the meaning of *In re Hoeksema*. In *Hoeksema*, the claims were rejected because an earlier patent disclosed a structural similarity to the Appellant's chemical compound. The court in *Hoeksema* stated: "We think it is sound law, consistent with the public policy underlying our patent law, that before any publication can amount to a statutory bar to the grant of a patent, its disclosure must be such that a skilled artisan could take its teachings in combination with his own knowledge of the particular art and be in possession of the invention." *In re Hoeksema*, 399 F.2d 269, 273, 158 USPQ 596, 600 (CCPA 1968). The meaning of *Hoeksema* for the present case is that

unless Leigh places Appellants' claims in the possession of a person of ordinary skill in the art, Leigh is legally insufficient to anticipate Appellants' claims under 35 U.S.C. § 102. As explained above, Leigh does not disclose each and every element and limitation of independent claims 1 and 8 of the present application. Because Leigh does not disclose each and every element, Leigh cannot possibly place the elements and limitations of the independent claims in the possession of a person of ordinary skill in the art. Leigh cannot, therefore, anticipate claims 1 and 8 of the present application.

Relations Among Claims

Independent claim 15 is a computer program product claim for determining scalability cabling between at least two scalable data processing systems corresponding to independent method claim 8. For the same reasons that Leigh does not disclose or enable a method for determining scalability cabling between at least two scalable data processing systems, Leigh also does not disclose or enable computer program products for determining scalability cabling between at least two scalable data processing systems corresponding to independent claim 15. Independent claim 15 is therefore patentable and should be allowed.

Claims 2, 9-10, and 16-17 depend from independent claims 1, 8, and 15. Each dependent claim includes all of the limitations of the independent claim from which it depends. Because Leigh does not disclose or enable each and every element of the independent claims, Leigh does not disclose or enable each and every element of the dependent claims of the present application. As such, claims 2, 9-10, and 16-17 are also patentable and should be allowed.

Argument Regarding The Second Ground Of Rejection On Appeal: Claims 3, 6-7, 11-12, 14, 18-19, and 21 Are Rejected Under 35 U.S.C. § 103(a) As Being Unpatentable Over Leigh In View Of Nakamura

Claims 3, 6-7, 11-12, 14, 18-19, and 21 stand rejected for obviousness under 35 U.S.C. § 103(a) as being unpatentable over Leigh in view of Nakamura, *et al.* (U.S. Publication

No. 2004/0057448) (hereafter, 'Nakamura'). The question of whether Appellants' claims are obvious or not is examined in light of: (1) the scope and content of the prior art; (2) the differences between the claimed invention and the prior art; (3) the level of ordinary skill in the art; and (4) any relevant secondary considerations, including commercial success, long felt but unsolved needs, and failure of others. *KSR Int'l Co. v. Teleflex Inc.*, No. 04-1350, slip op. at 2 (U.S. April 30, 2007). Although Appellants recognize that such an inquiry is an expansive and flexible one, the Office Action must nevertheless demonstrate a prima facie case of obviousness to reject Appellants' claims for obviousness under 35 U.S.C § 103(a). *In re Khan*, 441 F.3d 977, 985-86 (Fed. Cir. 2006). To establish a prima facie case of obviousness, the proposed combination of the references must teach or suggest all of the claim limitations of dependent claims 3, 6-7, 11-12, 14, 18-19, and 21. In re Royka, 490 F.2d 981, 985, 180 USPQ 580, 583 (CCPA 1974). Dependent claims 3, 6-7, 11-12, 14, 18-19, and 21 depend from independent claims 1, 8, and 15 and include all the limitations of the independent claims from which they depend. In rejecting dependent claims 3, 6-7, 11-12, 14, 18-19, and 21, the Office Action relies on Leigh as disclosing each and every element of independent claims 1, 8, and 15. As shown above, Leigh in fact does not disclose each and every element of independent claims 1, 8, and 15. Because Leigh does not disclose each and every element of independent claims 1, 8, and 15, the combination of Leigh and Nakamura cannot possibly disclose each and every element of dependent claims 3, 6-7, 11-12, 14, 18-19, and 21. The proposed combination of Leigh and Nakamura, therefore, cannot establish a prima facie case of obviousness, and the rejections 35 U.S.C. § 103(a) should be withdrawn.

Argument Regarding The Third Ground Of Rejection On Appeal: Claim 4 Is Rejected Under 35 U.S.C. § 103(a) As Being Unpatentable Over Leigh In View Of Nakamura and Vegter

Claim 4 stands rejected for obviousness under 35 U.S.C. § 103(a) as being unpatentable over Leigh and Nakamura further in view of Vegter (U.S. Patent No. 6,286,073) (hereafter, 'Vegter'). The question of whether Appellants' claim is obvious or not is examined in light of: (1) the scope and content of the prior art; (2) the differences

between the claimed invention and the prior art: (3) the level of ordinary skill in the art; and (4) any relevant secondary considerations, including commercial success, long felt but unsolved needs, and failure of others. *KSR Int'l Co. v. Teleflex Inc.*, No. 04-1350, slip op. at 2 (U.S. April 30, 2007). Although Appellants recognize that such an inquiry is an expansive and flexible one, the Office Action must nevertheless demonstrate a prima facie case of obviousness to reject Appellants' claims for obviousness under 35 U.S.C § 103(a). *In re Khan*, 441 F.3d 977, 985-86 (Fed. Cir. 2006). To establish a prima facie case of obviousness, the proposed combination of the references must teach or suggest all of the claim limitations of dependent claim 4. In re Royka, 490 F.2d 981, 985, 180 USPQ 580, 583 (CCPA 1974). Dependent claim 4 depends from independent claim 1 and includes all the limitations of the independent claim from which it depends. In rejecting dependent claim 4, the Office Action relies on Leigh as disclosing each and every element of independent claim 1. As shown above, Leigh in fact does not disclose each and every element of independent claim 1. Because Leigh does not disclose each and every element of independent claim 1, the combination of Leigh, Nakamura, and Vegter cannot possibly disclose each and every element of dependent claim 4. The proposed combination of Leigh, Nakamura, and Vegter, therefore, cannot establish a prima facie case of obviousness, and the rejections 35 U.S.C. § 103(a) should be withdrawn.

Argument Regarding The Fourth Ground Of Rejection On Appeal: Claims 5, 13, and 20 Are Rejected Under 35 U.S.C. § 103(a) As Being Unpatentable Over Leigh In View Of Vegter

Claims 5, 13, and 20 stand rejected for obviousness under 35 U.S.C. § 103(a) as being unpatentable over Leigh in view of Vegter. The question of whether Appellants' claims are obvious or not is examined in light of: (1) the scope and content of the prior art; (2) the differences between the claimed invention and the prior art; (3) the level of ordinary skill in the art; and (4) any relevant secondary considerations, including commercial success, long felt but unsolved needs, and failure of others. *KSR Int'l Co. v. Teleflex Inc.*, No. 04-1350, slip op. at 2 (U.S. April 30, 2007). Although Appellants recognize that such an inquiry is an expansive and flexible one, the Office Action must nevertheless

demonstrate a prima facie case of obviousness to reject Appellants' claims for obviousness under 35 U.S.C § 103(a). *In re Khan*, 441 F.3d 977, 985-86 (Fed. Cir. 2006). To establish a prima facie case of obviousness, the proposed combination of the references must teach or suggest all of the claim limitations of dependent claims 5, 13, and 20. In re Royka, 490 F.2d 981, 985, 180 USPQ 580, 583 (CCPA 1974). Dependent claims 5, 13, and 20 depend respectively from independent claims 1, 8, and 15, and include all the limitations of the independent claims from which they depend. In rejecting dependent claim 5, 13, and 20, the Office Action relies on Leigh as disclosing each and every element of independent claims 1, 8, and 15. As shown above, Leigh in fact does not disclose each and every element of independent claims 1, 8, and 15. Because Leigh does not disclose each and every element of independent claims 1, 8, and 15, the combination of Leigh and Vegter cannot possibly disclose each and every element of dependent claims 5, 13, and 20. The proposed combination of Leigh and Vegter, therefore, cannot establish a prima facie case of obviousness, and the rejections 35 U.S.C. § 103(a) should be withdrawn.

Conclusion of Appellants' Arguments

Claims 1-2, 8-10, and 15-17 stand rejected under 35 U.S.C. § 102 as being anticipated by Leigh. Leigh does not disclose each and every element of Appellants' claims and does not enable Appellants' claims. Leigh therefore does not anticipate Appellants' claims. Claims 1-2, 8-10, and 15-17 are therefore patentable and should be allowed. Appellants respectfully request reconsideration of claims 1-2, 8-10, and 15-17.

Claims 3, 6-7, 11-12, 14, 18-19, and 21 stand rejected under 35 U.S.C. § 103 as obvious over Leigh in view of Nakamura. The combination of Leigh and Nakamura does not teach or suggest each and every element of Appellants' claims. Claims 3, 6-7, 11-12, 14, 18-19, and 21 are therefore patentable and should be allowed. Appellants respectfully request reconsideration of claims 3, 6-7, 11-12, 14, 18-19, and 21.

Claim 4 stands rejected under 35 U.S.C. § 103 as obvious over Leigh in view of Nakamura and Vegter. The combination of Leigh, Nakamura, and Vegter does not teach or suggest each and every element of Appellants' claims. Claim 4 is therefore patentable and should be allowed. Appellants respectfully request reconsideration of claim 4.

Claims 5, 13, and 20 stand rejected under 35 U.S.C. § 103 as obvious over Leigh in view of Vegter. The combination of Leigh and Vegter does not teach or suggest each and every element of Appellants' claims.. Claims 5, 13, and 20 are therefore patentable and should be allowed. Appellants respectfully request reconsideration of claims 5, 13, and 20.

In view of the arguments above, reversal on all grounds of rejection is requested.

The Commissioner is hereby authorized to charge or credit Deposit Account No. 50-0563 for any fees required or overpaid.

Respectfully submitted,

Date: August 26, 2008 By:

H. Artoush Ohanian
Reg. No. 46,022
Biggers & Ohanian, LLP
P.O. Box 1469
Austin, Texas 78767-1469
Tel. (512) 472-9881
Fax (512) 472-9887
ATTORNEY FOR APPELLANTS

APPENDIX OF CLAIMS ON APPEAL IN PATENT APPLICATION OF ALFREDO ALDEREGUIA, *ET AL.*, SERIAL NO. 10/777,508

CLAIMS

Listing of Claims:

1. A scalable data processing system, including:

a first set of central processing units;

a first system memory accessible to the first set of processors;

scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system;

a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system; and

system management to cause each of the system's scalability ports to issue an identifiable signal and further configured to detect the reception of an identifiable signal, sent by another system, by any of the scalability ports and to report the reception of the signal to a system management of the second system to determine which ports of the two systems are connected by the cable.

2. The system of claim 1, wherein the system management includes a service processor connected to the system via an adapter card and wherein the service processor is connected to a service processor of other systems via a network medium.

3. The system of claim 1, wherein the system management causes a scalability port to issue an identifiable signal by causing the assertion of a bit in a register corresponding to the set of scalability ports.

4. The system of claim 3, wherein the scalability port register is implemented in a programmable logic device.

5. The system of claim 1, wherein the system management further includes controller logic connected to the service processor via a dedicated serial connection and connected to the programmable logic device via an I2C bus.

6. The system of claim 1, wherein the system management includes means for determining a timeout condition following assertion of a bit, and identifying the corresponding scalability port open.

7. The system of claim 1, wherein the system management further includes code means for using the scalability information to generate a graphical image of scalability interconections.

8. A method of determining scalability cabling between at least two scalable data processing systems, comprising:

driving an identifiable signal on a first scalability port of a first system;

responsive to receiving the identifiable signal by a second system, determining which scalability port of the second system received the distinctive signal;

informing the first system of the reception of the distinctive signal by the determined scalability port of the second system and recording the first scalability port of the first system and the scalability port of the second system as being connected by a scalability cable.

9. The method of claim 8, further comprising detecting a timeout by the first system and, responsive thereto, identifying the first scalability port as being unconnected.

10. The method of claim 8, further comprising, iterating the sequence of claim 8, until all scalability ports have been accounted for.

11. The method of claim 10, further comprising generating a graphical image of the scalability cable connections.

12. The method of claim 8, wherein driving the signal on the first scalability port comprises setting a bit in a register associated with the set of scalability ports.

13. The method of claim 8, wherein setting a bit in the register comprises a controller writing to a specified address on an I2C bus connecting the controller to the register.

14. The method of claim 8, wherein determining the scalability port that received the signal includes reading the bits in a register associated with scalability port.

15. A computer program product comprising computer executable instructions, stored on a computer readable medium, for determining scalability cabling between at least two scalable data processing systems, comprising:

computer code means for driving an identifiable signal on a first scalability port of a first system;

responsive to receiving the identifiable signal by a second system, computer code means for determining which scalability port of the second system received the distinctive signal;

computer code means for informing the first system of the reception of the distinctive signal by the determined scalability port of the second system and recording the first scalability port of the first system and the scalability port of the second system as being connected by a scalability cable.

16. The computer program product of claim 15, further comprising code means for detecting a timeout by the first system and, responsive thereto, identifying the first scalability port as being unconnected.

17. The computer program product of claim 15, further comprising, code means for iterating the sequence of claim 15, until all scalability ports have been accounted for.

18. The computer program product of claim 15, further comprising code means for generating a graphical image of the scalability cable connections.

19. The computer program product of claim 15, wherein the code means for driving the signal on the first scalability port comprises code means for setting a bit in a register associated with the set of scalability ports.

20. The computer program product of claim 15, wherein code means for setting a bit in the register comprises a controller writing to a specified address on an I2C bus connecting the controller to the register.

21. The computer program product of claim 15, wherein the code means for determining the scalability port that received the signal includes code means for reading the bits in a register associated with scalability port.

APPENDIX OF EVIDENCE ON APPEAL IN PATENT APPLICATION OF ALFREDO ALDEREGUIA, *ET AL.*, SERIAL NO. 10/777,508

This is an evidence appendix in accordance with 37 CFR § 41.37(c)(1)(ix).

There is in this case no evidence submitted pursuant to 37 CFR §§ 1.130, 1.131, or 1.132, nor is there in this case any other evidence entered by the examiner and relied upon by the Appellants.

RELATED PROCEEDINGS APPENDIX

This is a related proceedings appendix in accordance with 37 CFR § 41.37(c)(1)(x). There are no decisions rendered by a court or the Board in any proceeding identified pursuant to 37 CFR § 41.37(c)(1)(ii).